

SEU Mitigation Testing of Xilinx Virtex II FPGAs

C.C. Yui, G.M. Swift, *Member, IEEE*, C. Carmichael, R. Koga and J.S. George

Abstract—SRAM-based reconfigurable programmable logic is widely used in commercial applications and occasionally used in space flight applications because of its susceptibility to single-event upset (SEU). Upset detection and mitigation schemes have been tested on the Xilinx Virtex II X-2V1000 in heavy-ion and proton irradiation to control the accumulation of SEUs and to mitigate their effects on the intended operation. Non-intrusive upset detection and partial reconfiguration in combination with TMR can repair the design to maintain state information. In-beam results on a simple test design demonstrate the effectiveness of these methods when used together.

I. INTRODUCTION

The advancing microelectronic technologies applied to field-programmable-gate arrays (FPGAs) in the commercial sector has resulted in higher speed and lower core voltages, improving integration and allowing lower power consumption. The decreasing costs and development time needed to implement FPGAs compared to designs using discrete logic devices have made programmable logic devices favorable in space and avionic applications as well. They offer flexibility for changing requirements, in-system and on-orbit programmability as well as potential recovery of in-flight failures.

The Xilinx Virtex II is a re-configurable SRAM-based FPGA that also has the ability to conduct partial reconfiguration or, write to the configuration memory post-configuration without disturbing operation. However, while SRAM-based memory in the FPGA is useful for reconfiguration, the static memory elements and combinatorial logic paths are susceptible to upset from heavy-ion particles in interplanetary space. The Virtex II has been selected for the

present study because several variations of the Virtex FPGA are currently implemented in various missions as well as in several future missions. Several radiation studies have been carried out on SRAM-based FPGAs [1]-[3]. Extensive testing has also shown that, with proper mitigation, SEU induced failures can be properly controlled for the older Virtex family devices [4].

It is important to distinguish between SEU testing that is (a) static or (b) dynamic. Static testing involves quantifying upsets in the configuration memory elements without toggling clock, inputs and outputs of a fully-configured device during irradiation. Dynamic testing requires observation of a functional design under irradiation to determine the sensitivity of the combinatorial logic as well as upsets during transient signal propagation.

Static test results on the configuration memory of the Virtex II X-2V1000 along with projected upset rates have been reported by Yui, et al. [5]. These results are used as a comparison and baseline for data collected from recent dynamic tests. An on-going test effort is currently being conducted by the Xilinx Single-Event Effects Consortium, comprised of members from Xilinx, The Aerospace Corp., Sandia National Labs, SEAKR Engineering, Los Alamos National Lab, Information Sciences Institute, Jet Propulsion Laboratory and other currently joining members. The purpose of this effort is to study the single-event performance of complex Virtex II capabilities as well as various types of devices within the Virtex II family.

Following the static tests on the configuration memory [5], four dynamic test campaigns were conducted to understand the configuration memory, user and combinatorial logic in this test mode. The first two tests consist of dynamically monitoring configuration upsets during irradiation with heavy ions. The latter two tests consist of monitoring configuration, user and combinatorial logic and observing a functional design during irradiation while employing mitigation. The design configured into the DUT is eight shift registers clocking through a “checkerboard” pattern. Errors in the shift register output and configuration memory read operations were used to gauge the system’s performance.

Manuscript received July 31, 2003. The research done in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration (NASA), under the NASA Electronic Parts and Packaging Program (NEPP), Code AE and in collaboration with Xilinx Inc. and The Aerospace Corp.

C.Y. and G.S. are with the Jet Propulsion Laboratory, Pasadena, CA 91109 USA (telephone: 818-393-0945, 818-354-5059, e-mail: candice.yui@jpl.nasa.gov, gary.m.swift@jpl.nasa.gov).

C.C. is with Xilinx Inc., San Jose, CA 95124 USA (telephone: 408-879-5114, e-mail: carl.carmichael@xilinx.com).

R.K. and J.G. are with The Aerospace Co., Los Angeles, CA 90009 (telephone: 310-336-6583, 310-336-2089, e-mail: rokutaro.koga@aero.org, jeffrey.s.george@aero.org).

II. EXPERIMENTAL DETAILS

A. Device Information

The device chosen for this study was the Virtex II X-2V1000. It was fabricated using the QPro radiation evaluation sample mask set for the XQR2V1000 but unlike the XQR-prefix device, it was produced on bulk CMOS wafers (that is, without an epitaxial layer). The mask is identical to one intended for XQR (or Xilinx QPro) line of radiation hardened Virtex devices. The absence of the epitaxial layer is unlikely to change upset cross-sections very much and is used in the XQR-line to eliminate single-event latch up. (A single-event latch up test was done separately on the XQR2V3000 and no latch ups occurred for LET up to 104 MeV-cm²/mg.)

The device is packaged in a 256-pin wire-bond standard ball gate array (BGA) package. It is fabricated on a 0.15μm CMOS 8-layer metal process and includes 40 block RAMs (737,280 total bits), 432 maximum I/Os, and 2.8M configuration bits. These devices were obtained for the sole purpose of SEU testing. The X-2V1000 was ideal for SEU characterization because it is one of two members of the Virtex II family that has a face-up die, suitable for heavy ion penetration after de-lidding. Prior to testing, the device was chemically etched to expose the die. Testing was done at the Cyclotron Institute, Texas A&M University; Lawrence Berkeley National Laboratory 88" cyclotron, UC Berkeley and at the Crocker Nuclear Laboratory, UC Davis. Devices were tested at normal incidence for an LET range of 1.5 – 63 MeV/ (cm²/mg) and proton energy of 6.8 MeV. Testing was conducted in a vacuum chamber at the Lawrence Berkeley test facility and in-air at the other facilities.

B. Test Setup

The test platform for the dynamic testing is shown in Fig. 1 and 2. It is composed of three primary components, the device-under-test (DUT), the Configuration Monitor and the Functional Monitor. The DUT is situated on a Xilinx HW-AFXBG256-200 development board incorporating a Xilinx XCV100 Virtex FPGA. The XCV100, also known as “service FPGA” acts as the “Configuration Monitor”, detecting and correcting configuration upsets under the control of a host computer. The service FPGA continuously detects and counts the number of errors in the DUT through the use of readback, a feature of Xilinx FPGAs that allow users to read the configuration memory post-configuration. The FPGA readback data from the configuration memory undergoes a bit-for-bit comparison with a mask file stored in a separate PROM, also known as “mask PROM”. Mismatches in the bit-for-bit comparison are classified as detected errors and a pulse is sent from the XCV100 to a host computer via custom Visual Basic software and 40 pin twisted-pair ribbon cables.

Once errors are detected, the service FPGA also has the ability to correct the upsets through partial re-configuration (PRC). This process, also known as “non-intrusive scrubbing”, will cause errors in the configuration memory to be corrected without interfering with the operation of the loaded design [6]. The custom Visual Basic program records and displays the errors as they occur. It also gives the user the option to continuously monitor, continuously scrub, or continuously monitor and scrub at 50% duty cycles.

The “Functional Monitor” is implemented using another FPGA, a Spartan XCS30XL, used to generate test vectors to the DUT and compare DUT outputs with expected values.

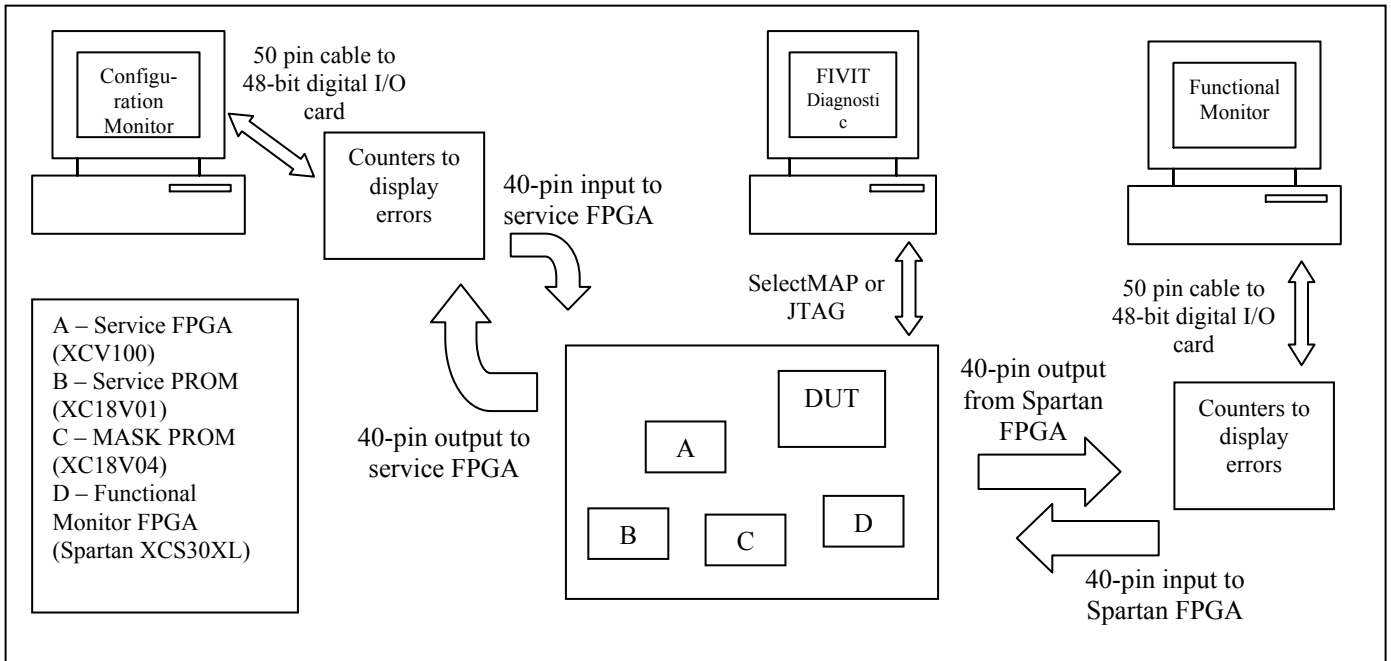


FIGURE 1




Fig. 2. A photo of the test vehicle in the vacuum chamber at Lawrence Berkeley National Laboratory 88" cyclotron.

Errors in the comparison are then sent via ribbon cables to a second host computer and recorded by a different Visual Basic program. All three FPGAs are configured using Xilinx one-time programmable or in-system programmable PROMs.

Counters with 9-digit LED displays on both the Configuration Monitor and Functional Monitor allowed the user to visually see configuration and functional errors as they occurred. A custom C++ software application was also available at the end of each beam run to read back the number of errors that accumulated in the configuration logic block (CLB) frames, block RAM cells and configuration control registers. The custom software is named FIVIT for Fault Injection Verification Tool and communicates with the DUT via a SelectMAP or JTAG parallel III interface. Only the JTAG interface was used in our dynamic testing. FIVIT was helpful in verifying the efficacy of the Configuration Monitor and its ability to detect and fix configuration upsets. The first two dynamic tests only monitored and correct errors and did not include the Functional Monitor. Power was supplied to the test platform from an HP 6629A digital power supply and recorded by a separate computer and strip-charting program. An Agilent function generator was used to supply a 3.3V square wave to the clock input of the DUT.

C. Test Design

The DUT design consisted of eight simple shift registers, made of 500 flip-flops each, using roughly 40% of the available flip-flops. A second design identical to the previous design was also tested with the exception that four of the eight shift registers implemented triple module redundancy (TMR), a technique that will be further explained below. This design used about 80% of the available flip-flops. Inputs to the shift registers are provided by the Functional Monitor FPGA where a pattern of all zeroes, ones, or checkerboard can be selected

by the user. The clock frequency used for these tests was 10 kHz.

Future testing will be needed to analyze the use of more complicated resources in the Virtex II such as look-up tables, multiplier blocks, and digital clock managers and their sensitivity to SEU. The amount and types of resources used are highly influential on the sensitivity of a non-mitigated design.

III. TEST RESULTS

A. Dynamic Testing

To harden the Virtex II for space applications, a design using triple module redundancy and partial reconfiguration was tested. Triple module redundancy is a design technique widely used since both user logic and logic paths are susceptible to SEUs, a fact among reconfigurable logic devices. By implementing three full copies of the base design in the FPGA, SEUs and single-event transients (SETs) are voted out as they occur in the user or combinatorial logic.

The use of partial reconfiguration [6] in the Virtex and Virtex II devices repairs errors in the configuration bitstream without disrupting the operation of the device. Scrubbing non-intrusively allows individual configuration bits to be changed (back to their correct value following an upset) in individually addressable "frames". Prior to scrubbing, the user can choose whether to require a readback operation to determine if errors exist. A detailed picture illustrating the configuration method is found in Fig. 3.

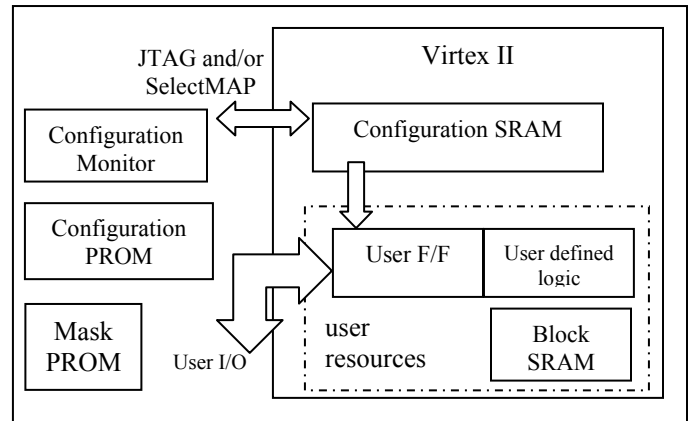


Fig. 3. A diagram of the device resources involved in configuration and partial reconfiguration. The configuration PROM contains the bitstream used to program the configuration monitor FPGA. The configuration monitor FPGA in turn, monitors and corrects errors as seen in the DUT's SRAM-based configuration memory through either the JTAG or SelectMAP port. The DUT's configuration dictates how the design will be implemented using the user resources. The mask PROM is used to flag the bits in the configuration bitstream that should not be changed by the Configuration Monitor.

Eliminating configuration errors as they occur prevents the accumulation of errors in which functional failure of the

device is a likely consequence. In addition, upsets in the configuration memory creates driver contention where two inverter outputs of different states can be connected and cause a local abnormally high current. Initial configuration and scrubbing of the device is made possible through the JTAG or SelectMAP ports.

The only effective mitigation method for an SRAM-based FPGA design intended for space flight is to incorporate both TMR and scrubbing. When combined and used appropriately, single errors in the user or path logic and static errors in the memory can be corrected before the next error occurs.

Before comparing mitigated test results, it is interesting to compare configuration upsets in static and dynamic tests. Fig. 4 shows three cross-section curves for average configuration memory and block RAM upsets. Pure static test results from a simple single shift-register design are shown (solid diamonds). These are very close to the dynamic results where the configuration was being monitored continuously (the open square points). However, dynamic monitoring of the somewhat more complex eight-shift register (four with TMR) gives a significantly higher cross-section. This suggests a design dependence that is greater than the difference between static and dynamic testing (at least where dynamic is defined as continuously reading the configuration memory).




Fig. 4. The cross-section of combined configuration memory and block RAM bits for a static test and dynamic test without mitigation efforts.

TMR only, partial reconfiguration only and combined TMR and scrub designs were tested in protons and heavy ions. In each case, the user logic upsets as well as functional failure of the shift registers were strip-charted and time-stamped by the custom software during the irradiation. A functional failure is defined here to be the state where the shift register loses its ability to shift the data out correctly and is determined by a signature of a constant stream of errors. The fluence to upset can be obtained by calculating the fluence accumulated at the point of failure. When scrubbing was enabled, it was important to make certain the upset rate was less than the scrub rate. Overwhelming the test system with more upsets than it is designed to mitigate would produce misleading and

erroneous data. The scrub frequency for the 2.8M configuration and 737,280 block RAM bits (3.5M total bits) used for this test was 20 MHz, the maximum frequency of our communication hardware and software. However, this scrub frequency was reduced by half due to the fact that the test mode used most often was alternating readback and scrub cycles. Therefore, in order to maintain conditions suitable for the scrub rate, the maximum upset rate is 3 upsets/second. This provided a challenge during testing as the flux, fluence and experimental results had to be carefully balanced to guarantee meaningful test data.

Before testing with heavy ions, the test vehicle was brought to the Crocker Nuclear Laboratory for proton testing to verify its functionality as well as to obtain preliminary results to compare with previous non-mitigated data. TMR and partial reconfiguration mitigation methods were tested at 6.8 MeV. A scatter plot of the fluence to first functional failure is shown in Figure 5. A slightly higher fluence is observed for runs employing partial reconfiguration and one run for TMR. The highest fluence level for a non-mitigated design is about twice than the highest fluence for other designs. No functional errors were recorded for the design using both TMR and partial reconfiguration.




Fig. 5. A scatter plot of the fluence to first functional failure for protons. "PRC" represents data points where partial reconfiguration was used, "TMR" represents points where triple module redundancy was used in the design.

A summary of the test results from the heavy ion test can be seen in Figure 6. Five test runs with low upset rates from each mitigation category are compared. When mitigation such as TMR is used in the design, a decrease of roughly 25% in frequency of functional errors is seen. When partial reconfiguration is implemented, functional errors decrease by about 40%. When both mitigation methods are used in conjunction to repair configuration memory upsets as well as user logic upsets, no functional errors in the shift register were observed. However, it should be noted that a single bit error was seen in one shift register output during one of the beam runs where TMR and PRC was used. Some of the possible

causes of this event could be an unscrubbed single-event transient or an ion strike to the enable pin of the TMR minority voter [7]. Although the probability of seeing an error just prior to or following a scrub is also not impossible, a mitigated design is useful in controlling the upsets as they occur.




Fig. 6. A comparison of frequency of errors to total runs for four possible variations of the same shift register design shows the effectiveness of each method.

A comparison of cross-section for first functional errors is given in Fig. 7. The varied scattering seen before in Fig. 5 is replicated here and is a function of Poisson's probability distribution, a mathematical computation of the probability of atypical events occurring in a given time period. Overall, a majority of the cross-sections for the non-mitigated design are slightly higher than all points for a design using TMR and a few points using scrubbing. The present data is constrained by limited statistics and would be helped by additional testing.




Fig. 7. Cross-section of first functional errors in heavy-ions. Poisson scatter is large due to the small number of runs at each LET.

From the observations made in the recent dynamic tests, using partial reconfiguration and triple module redundancy is beneficial for space applications using the Virtex II. However, one thing to consider when using TMR is the fact

that the user design is constrained to using one-third of the device functionality. Sufficient resources must be left available to implement three total copies of the user design.

To apply partial reconfiguration, a controller and an extra PROM to store the configuration bitstream mask file is needed. The controller can be a reprogrammable or one-time programmable FPGA. The role of the controller would be similar to the previously described "Configuration Monitor," constantly detecting and repairing configuration memory errors.

B. Single-Event-Functional Interrupts (SEFIs)

Besides upsets to the configuration, user and path logic, there are also three noted single-event-functional-interrupts (SEFIs) in which the DUT is susceptible. They include the power-on-reset (POR) SEFI, SelectMAP SEFI and a JTAG SEFI and are results of ion hits to the corresponding circuitry. Their effects, detection and recovery mechanisms are detailed in Table 1. Typically, these SEFIs are low in occurrence [5] and would rarely occur while in orbit. However, in test environments where conditions are generally a worse-case example of the intended use of the test device, SEFIs can be examined. The criterion for a SEFI for the Virtex II is that it requires an intrusive full reset and complete reconfiguration of the device before returning to normal operability.

TABLE I
SUMMARY OF SEFI GROUPS

POR	
Effects	Clearing of configuration memory and loss of state data
Detection	Done pin transitions low, I/O becomes tri-stated, no user functionality available
Recovery	Standard configuration. No power cycle necessary.
SelectMAP	
Effects	Loss of communication with configuration logic. Configuration error detection and non-evasive correction unavailable.
Detection	Not responsive to data readbacks
Recovery	Standard configuration. No power cycle necessary.
Effects	Loss of communication with configuration logic. Configuration error detection and non-evasive correction unavailable.
Detection	Read access to configuration memory returns constant value.
Recovery	Standard configuration. No power cycle necessary.

Note: Standard configuration is intrusive, that is, FPGA user design is temporarily disabled during reloading of configuration memory and reset of control logic.

The dynamic test vehicle had two POR SEFI detection mechanisms, one in the Configuration Monitor service FPGA and one in the Functional Monitor. The service FPGA continuously monitors the state of the "DONE" pin of the device for a transition to low, a sign that the POR SEFI has occurred. When this occurs, the service FPGA will automatically reconfigure the DUT and increase the

number in an internal SEFI counter. The “DONE” signal was also routed to the Functional Monitor for observation along with two feed-through signals added to the DUT design. Pin assignments for these two signals were chosen to be in close proximity to each other. An unexpected state in any of these feed-throughs would indicate that the configuration memory was likely to be severely corrupted and has lost the ability to maintain its configuration.

Corrupted bits in the frame address register (FAR) were also continuously watched. The FAR register contains the address in the configuration data frames where partial reconfiguration will begin for either the SelectMAP or JTAG port. It is recommended that FAR functionality is checked prior to reconfiguration. The consequence of an error in the FAR is similar to a SEFI and for this reason, it is also classified as a SEFI. However, its incidence has been extremely low and has been negligible compared to other SEFIs.

No detection capability was implemented for the SelectMAP SEFI as only the JTAG was used during the experiments. The JTAG SEFI is a loss of communication with the configuration logic and is detected by a constant value being returned by the configuration memory read operation.

Two individual runs with a SEFI event were recorded during the heavy ion dynamic testing. The first event was a JTAG SEFI where the JTAG could no longer read to the configuration memory or perform a write operation, demonstrated by an inability to scrub. The second event was a POR SEFI where a transition in the “DONE” signal was observed and functionality in all eight shift registers was lost at the time of a sudden large increase in configuration memory upsets.

The mechanism of SEFIs is independent of the mitigation involved and is inherent in the device architecture. However, when proper mitigation is used to remove all possibilities of errors from single-event upsets within a single device, applying FPGA device redundancy can potentially remove the consequence of SEFIs, thus producing a robust system for critical flight applications.

IV. CONCLUSIONS

In this study we have demonstrated the effectiveness of the mitigation techniques of TMR and PRC when used in combination for the Virtex II X-2V1000. A comparison of the frequency of functional failures shows that the designs utilizing mitigation techniques such as partial reconfiguration or TMR alone have only a slight advantage over a non-mitigated design. However, when both methods are used, the design was observed to be essentially immune to functional errors. Supporting results are also demonstrated by the fluence and cross-section data of first functional errors. These initial results on a simple test design are encouraging and

suggest that using TMR and partial reconfiguration mitigation methods together can make the Virtex II suitable for space flight applications. More testing is needed on designs of greater complexity to confirm that this result is applicable and comparable with the present results.

V. ACKNOWLEDGMENT

The authors wish to thank the Xilinx Consortium for their support. Special thanks are also given to Travis Minto, Duc Nguyen and Michael O'Connor of the Jet Propulsion Laboratory for their preparation of the test hardware.

VI. REFERENCES

- [1] E. Fuller, M. Caffrey, P. Blain, C. Carmichael, N. Khalsa, and A. Salazar, “Radiation test results of the Virtex FPGA and ZBT SRAM for Space Based Reconfigurable Computing,” 1999 MAPLD.
- [2] J.J. Wang, R.B. Katz, J.S. Sun, B.E. Cronquist, T.M. Speers, and W.C. Plants, “SRAM Based Re-programmable FPGA for Space Applications,” IEEE Trans. Nucl. Sci., vol. 46, no. 6, 1728-1735, Dec. 1999.
- [3] D. Brodrick, A. Dawood, N. Bergmann, M. Wark, “Error Detection for Adaptive Computing Architectures in Spacecraft Applications,” 6th Australasian Computer Systems Architecture Conference, 2001.
- [4] C. Carmichael, E. Fuller, P. Blain, and M. Caffrey, “SEU Mitigation Techniques for Virtex FPGAs in Space Applications,” MAPLD, 1999.
- [5] C. Yui, G. Swift, and C. Carmichael, “Single-Event Upset Susceptibility of the Xilinx Virtex II FPGA,” MAPLD, 2002.
- [6] C. Carmichael, “Correcting Single-Event Upsets Through Virtex Partial Configuration,” Xilinx Application Note XAPP216, June, 2000.
- [7] C. Carmichael, “Triple Module Redundancy Design Techniques for Virtex FPGAs,” Xilinx Application Note XAPP197, November, 2001.